

ORIGINAL RESEARCH

Open access

A Human-in-the-Loop Optimization Model for Demand Response Dispatch Under Comfort Constraints, Prosumer Preferences, and Low-Carbon Flexibility Signals

Charlotte Dupuis^{1*}, Hugo Perrin¹

Abstract

Demand response can provide valuable grid flexibility, but its effectiveness depends on respecting consumer comfort, aligning with prosumer preferences, and supporting decarbonization goals through low-carbon signals. Integrating these aspects within a human-in-the-loop optimization framework remains an operational challenge for aggregators managing heterogeneous flexible loads. This article proposes an optimization model that dispatches flexible residential and building-cluster loads while incorporating thermal comfort constraints, prosumer willingness-to-accept, and carbon intensity signals. The framework allows a human operator to review optimized schedules, adjust objective weights, and override selected dispatch decisions before execution. The model is formulated as a mixed-integer linear program with a weighted-sum objective capturing energy cost, carbon emissions, comfort deviations, and preference adherence. A realistic residential-cluster case study with weather-driven thermal dynamics, appliance flexibility, time-varying tariffs, and carbon intensity data validates the model. The proposed dispatch reduces feeder peak demand by 18.6% and carbon emissions by 11.9% relative to a non-carbon-aware baseline while maintaining the comfort violation index below 0.035 degree-hours per occupied interval. Operator overrides occur in 6.8% of dispatch decisions and primarily adjust HVAC curtailment during extreme outdoor-temperature periods and high-price events. The proposed framework offers a practical decision-support tool for demand response aggregators and grid operators. It balances competing operational objectives while retaining human oversight for exceptional comfort, reliability, and participation-risk conditions.

Keywords Demand response, Human-in-the-loop optimization, Comfort constraints, Prosumer preferences, Low-carbon flexibility, Dispatch

*Correspondence:

Charlotte Dupuis
charlotte.dupuis@gmail.com

¹ Department of Smart Energy Networks, Faculty of Engineering, University of Lille, Lille, France

Introduction

Demand response has evolved from simple price-based load shifting into an operational flexibility resource capable of supporting distribution networks, renewable integration, and peak reduction. Residential and commercial flexible loads such as HVAC systems, electric water heaters, and deferrable appliances can be coordinated through optimization, but purely cost-driven scheduling may create rebound peaks or unacceptable service disruptions [1]. Recent control studies show that building thermal inertia and predictive scheduling can convert HVAC loads into dispatchable resources when indoor comfort dynamics are modeled explicitly [2]. This article addresses the resulting coordination problem by

formulating a human-in-the-loop demand response optimizer that jointly values grid flexibility, comfort preservation, prosumer preferences, and carbon-aware operation.

Comfort is not a secondary consideration in demand response, because thermal discomfort and unwanted appliance interruptions directly affect participation durability. Occupant-centric studies emphasize that HVAC control must account for personal comfort, override behavior, and zone-level heterogeneity rather than assuming a uniform temperature set point [3]. Empirical evidence from thermostat override studies shows that users actively intervene during demand response events when automation conflicts with perceived comfort or convenience [4]. Therefore, the proposed model embeds comfort bounds and soft violation variables directly into the optimization problem rather than evaluating comfort only after dispatch.

Prosumer flexibility is also shaped by willingness-to-accept, distrust of direct load control, and perceived fairness of compensation. Survey and behavioral studies indicate that enrollment in residential demand response depends on both monetary incentives and non-monetary factors such as control autonomy, technology trust, and the type of load being controlled [5]. Broader reviews of consumer engagement similarly show that sustained participation requires transparent value propositions and sensitivity to household routines [6]. The model developed here represents this behavioral dimension through segment-specific flexibility functions, appliance-level participation limits, and penalty terms for deviations from stated prosumer preferences.

Low-carbon flexibility signals add a further operational layer by encouraging load increases during cleaner generation periods and curtailment during carbon-intensive intervals. Carbon-aware demand response has been examined through marginal emission factors, time-varying carbon signals, and integrated energy-system scheduling, but these signals can conflict with comfort and price objectives [7]. Building-cluster and multi-energy studies demonstrate that carbon-oriented dispatch can reduce emissions when the carbon signal is mapped into either objective weights or explicit operating constraints [8]. The contribution of this article is a mathematically explicit MILP that combines these elements with operator review, enabling human adjustment of trade-off weights and selective overrides before the optimized schedule is implemented.

Demand Response Problem

Types of demand response programs

The dispatch problem considers three program classes: price-based demand response, incentive-based direct load control, and carbon-aware flexibility activation. Price-based programs shift appliance operation away from high-tariff intervals, as in residential load-shifting formulations using heuristic and metaheuristic optimization [9]. Incentive-based programs allow an aggregator to request curtailment or rescheduling from enrolled prosumers, with compensation linked to inconvenience and flexibility depth [10]. Carbon-aware programs extend the dispatch signal by using time-varying emissions information so that flexible consumption is moved toward lower-carbon periods where technically and behaviorally feasible [11].

Dispatch objectives and constraints

The operational objective is to minimize a weighted sum of energy procurement cost, peak demand penalty, carbon emissions, comfort deviation, and prosumer preference violation. HVAC loads are represented through first-order thermal dynamics, water heaters through temperature-state constraints, and appliances through binary start-time variables, following established demand response models for thermostatically controlled and deferrable loads [12]. Feeder-level constraints impose transformer capacity and ramping limits, while household-level constraints preserve service completion and maximum interruption duration. This creates a mixed-integer scheduling problem whose feasible region is shaped by both electrical network requirements and user-centered comfort boundaries.

The human-in-the-loop perspective

A human-in-the-loop architecture is introduced because automated optimization may recommend technically optimal dispatch schedules that are operationally undesirable under unusual weather, reliability warnings, or participation-risk conditions. Prior human-centered building-control research shows that occupant feedback and operator interpretation can improve energy decisions when automation cannot fully infer comfort or social acceptability [13]. In the proposed framework, the optimization engine produces a candidate schedule, after which the operator reviews predicted peak reduction, comfort deviation, household participation burden, and carbon savings. The operator may then adjust objective weights or lock selected load decisions, and the MILP is re-solved with these human-imposed constraints before final dispatch.

Figure 1 presents the proposed human-in-the-loop demand response architecture linking household flexibility inputs, comfort and preference constraints, carbon-aware optimization, operator review, re-optimization, and final dispatch execution.




Figure 1. Human-in-the-Loop Demand Response Optimization Architecture for Comfort-Aware and Low-Carbon Flexibility Dispatch

Comfort Constraint Modeling

Thermal comfort models

Thermal comfort is modeled using indoor temperature deadbands, occupied-period penalties, and optional predicted mean vote approximations where building metadata are available. Occupant-centric HVAC studies show that personalized comfort models can improve control performance compared with fixed set-point assumptions, especially when households differ in preferred temperatures and tolerance ranges [14]. For tractability, this article uses linear comfort envelopes around preferred temperature trajectories rather than nonlinear full-body heat-balance equations. The resulting comfort formulation is compatible with MILP dispatch while retaining the operational meaning of temperature acceptability.

Occupancy and behavioral variability

Occupancy affects both the severity and timing of comfort constraints because deviations during occupied evening periods are more disruptive than deviations during unoccupied daytime intervals. Demand response experiments in buildings show that occupant-oriented control must distinguish between zones, schedules, and behavioral responses to curtailment events [15]. Thermostat override evidence further indicates that households respond asymmetrically to heating and cooling discomfort, making a single uniform penalty insufficient [16]. The proposed model therefore assigns

time-dependent comfort weights, household-specific override-risk coefficients, and appliance-specific acceptability windows.

Comfort constraint formulation

For each dwelling (i) and interval (t), indoor temperature $T_{i,t}$ must remain within a preferred lower and upper comfort

$$T_{i,t} \geq T_{i,t}^-$$
$$T_{i,t} \leq T_{i,t}^+$$

bound unless a nonnegative slack variable is activated. The linear constraints are $(-s_{i,t}^- \text{ and } T_{i,t} \leq T_{i,t}^+ + s_{i,t}^+)$, where $s_{i,t}^-$ and $s_{i,t}^+$

measure heating and cooling discomfort. Similar soft-constraint structures have been used in predictive demand response control to balance thermal inertia, grid signals, and occupant satisfaction [2]. The objective penalizes occupied-period slack more heavily than unoccupied-period slack, thereby allowing limited flexibility without treating comfort as infinitely rigid.

Parameter selection

Comfort parameter values are selected from building-control literature and calibrated to the case-study climate using observed outdoor temperature, baseline thermostat schedules, and load response. Heat-pump and HVAC demand response studies indicate that wider deadbands increase flexibility but can increase dissatisfaction if curtailment duration is not limited [17]. Direct load control participation studies also suggest that consumers are more willing to offer flexibility when discomfort is bounded, transparent, and compensated [18]. Table 1 summarizes the comfort parameters used in the numerical study and defines how each load class contributes to the optimization.

Table 1. Comfort and Service Constraints for Load Classes in the Numerical Optimization Study

Load type	Comfort or service constraint	Parameter value used in study	Time dependence	MILP representation
HVAC cooling	Occupied indoor temperature band	22.5–25.5 °C	17:00–23:00 high penalty; other occupied hours medium penalty	Linear thermal state equation with upper/lower soft bounds
HVAC heating	Occupied indoor temperature band	20.0–22.5 °C	Morning and evening high penalty	Linear thermal state equation with asymmetric slack penalties
Electric water heater	Minimum tank temperature	48 °C normal, 45 °C emergency flexibility floor	Morning and evening draw periods high penalty	Linear temperature-state balance with minimum-temperature slack
Dishwasher	Latest completion time	Within 6 hours after user availability window opens	Evening start windows preferred	Binary start-time variable with completion constraint
Clothes washer	Latest completion time	Within 8 hours after declared availability	Weekend penalty reduced by 25%	Binary start-time variable and inconvenience penalty
EV charging, optional prosumer resource	Required departure energy	90% of declared energy requirement by 07:00	Overnight flexibility only	Energy balance with binary charging status and departure constraint

Prosumer Preference Modeling

Willingness-to-accept and flexibility functions

Prosumer preferences are represented through willingness-to-accept functions that map load type, delay duration, curtailment depth, and time of day into an inconvenience cost. Empirical studies of residential flexibility show that willingness to enroll differs sharply across household loads, with direct control of comfort-related devices requiring stronger compensation than deferrable appliance scheduling [10]. The model assigns each prosumer (i) a flexibility coefficient (α_i) and each appliance (a) a delay sensitivity (β_a), producing an inconvenience term proportional to accepted shift duration and service criticality. This structure allows the optimizer to select households with lower marginal inconvenience before calling on more reluctant participants.

Preference elicitation and segmentation

Preference parameters are obtained through a two-stage process combining enrollment questionnaires with revealed responses from past demand response events. Survey-based evidence shows that residential users can be segmented by automation acceptance, environmental motivation, and sensitivity to financial incentives [18]. Agent-based demand response enrollment studies further demonstrate that participation rates and response quality vary across consumer classes, motivating segment-specific parameterization rather than a single representative household [19]. In the case study, prosumers are grouped into comfort-prioritizing, cost-responsive, carbon-motivated, and automation-friendly segments, each with distinct participation limits and inconvenience coefficients.

Integration into optimization

Preferences enter the MILP through both constraints and objective penalties. A household-specific participation constraint limits the number of dispatch events over a rolling horizon, while an appliance-level preference penalty discourages schedules that violate stated availability windows or maximum delays. Peer-to-peer and prosumer-centered energy management studies show that preference-driven optimization can preserve autonomy while still achieving system-level coordination [20]. The proposed formulation extends that idea to aggregator dispatch by allowing the operator to observe the marginal inconvenience cost of each selected flexibility action before approval.

Sensitivity to preference uncertainty

Preference uncertainty is modeled by perturbing willingness-to-accept parameters and participation availability in the sensitivity analysis. Consumer distrust and perceived loss of control can reduce willingness to participate in direct load control even when compensation is offered, which makes deterministic preference assumptions operationally fragile [5]. Engagement reviews also show that nonresponse and opt-out behavior are common barriers when demand response programs fail to match household expectations [6]. The model therefore includes reserve flexibility margins and re-optimization after operator overrides, reducing dependence on any single uncertain participation estimate.

Low-Carbon Flexibility Signals

Sources of carbon intensity data

The low-carbon signal is defined as a time-indexed estimate of grid carbon intensity in grams of carbon dioxide equivalent per kilowatt-hour. Carbon-aware demand response studies use either average grid carbon intensity, marginal emission factors, or forecasted time-varying emission factors depending on data availability and decision purpose [7]. For short-term dispatch, marginal emission factors are operationally preferred because they approximate the emissions effect of incremental load shifting more directly than annual average factors [8]. The case study uses an hourly forecast signal updated daily and mapped to the 15-minute dispatch resolution by piecewise-constant interpolation.

Mapping carbon signals to dispatch

Carbon signals are mapped into dispatch by assigning each load action an emissions coefficient equal to the product of scheduled consumption and interval carbon intensity. Multi-energy demand response research shows that time-varying carbon factors can be integrated with price and incentive signals to guide load toward cleaner periods [11]. For flexible consumption that can be delayed rather than curtailed, the optimization compares the carbon intensity of the original service interval with feasible alternative intervals. This enables the model to reward load shifting that reduces emissions while penalizing shifts that merely move demand from a low-carbon period into a high-carbon one.

Carbon objective formulation

The carbon component of the objective is written as $(\lambda_C \sum_t \sum_i \sum_a, \gamma_t p_{i,a,t} \Delta t)$, where (γ_t) is carbon intensity, $(p_{i,a,t})$ is scheduled power, and (λ_C) converts emissions into the same normalized objective scale as cost and comfort. Carbon-oriented demand response models for data centers and building clusters demonstrate that emissions reductions depend strongly on signal quality and temporal flexibility windows [21]. To preserve tractability, this article uses a linear carbon term rather than nonlinear lifecycle accounting. The operator can increase (λ_C) during decarbonization events or reduce it during reliability emergencies when peak reduction is more urgent.

Trade-offs with cost and comfort

Cost, comfort, and carbon objectives do not always align because low-price periods may coincide with high-emission generation, and low-carbon periods may occur during occupied hours when flexibility is limited. Multi-objective economic dispatch studies show that demand response can reduce system cost while producing different emissions outcomes depending on renewable obligations and weighting choices [22]. Low-carbon demand response formulations under uncertainty further indicate that robust treatment of carbon signals may reduce emissions but can increase dispatch conservatism [23]. Table 2 defines the carbon-signal inputs used in the case study and specifies how each signal enters the MILP and operator review screen.

Table 2. Carbon-Signal Inputs for MILP Optimization and Operator Review in the Case Study

Signal element	Data source type used in study	Temporal granularity	Threshold or transformation	Optimization role	Operator display role
Forecast average carbon intensity	Day-ahead grid carbon forecast	1 hour mapped to 15 minutes	Low: <180 gCO ₂ e/kWh; medium: 180–350; high: >350	Linear emissions term in objective	Shows expected emissions profile by dispatch interval
Forecast marginal emission factor	Estimated marginal generator signal	1 hour	Normalized to 0–1 dispatch score	Sensitivity case replacing average signal	Flags intervals where incremental load is most carbon-intensive
Carbon event indicator	Aggregator-defined high-carbon alert	15 minutes	Binary value activated above 350 gCO ₂ e/kWh	Optional constraint limiting discretionary load	Highlights periods requiring operator confirmation
Carbon price equivalent	Scenario parameter	Daily value applied across horizon	€0, €75, and €150 per tCO ₂ e	Converts emissions into monetary-equivalent objective term	Allows operator adjustment of carbon-cost weight

Clean-flexibility opportunity index	Derived from carbon difference between feasible start times	15 minutes	Difference between baseline and optimized emissions	Rewards beneficial shifting within service windows	Ranks loads by emissions-saving potential
Carbon-comfort conflict score	Derived from carbon intensity and occupied comfort penalty	15 minutes	Product of normalized carbon score and comfort weight	Not a hard constraint; used in override review	Identifies cases where low-carbon action may create discomfort

Optimization Formulation

Decision variables and sets

The optimization model is defined for prosumers, flexible devices, and dispatch intervals, with device schedules represented through operational status, scheduled power, comfort state, and dispatch eligibility variables. HVAC systems, water heaters, appliances, and optional charging loads are modeled according to their technical flexibility and service requirements, consistent with multi-district and building-cluster demand response formulations [24]. The formulation distinguishes between controllable flexibility and unavailable load so that the aggregator cannot dispatch devices outside declared participation windows. Human-reviewed decisions are represented by schedule-locking rules that preserve operator-approved changes during re-optimization.

Objective function

The objective minimizes the combined operational burden of energy procurement, feeder peak exposure, carbon impact, comfort deviation, and preference-regret cost. This structure follows prosumer-centered optimization studies in which flexibility is treated as a heterogeneous and preference-sensitive resource rather than a uniform technical commodity [20]. The carbon component is included as a weighted dispatch signal, while the preference component discourages repeated or inconvenient control of the same household. The operator can adjust the relative importance of these components before execution, allowing the optimization to reflect system priorities under normal, constrained, or low-carbon operating conditions.

Constraints: grid, comfort, preference, and carbon

Grid constraints ensure that aggregate feeder loading remains within operational limits and that dispatch actions do not create excessive rebound after an event. Comfort constraints maintain indoor temperature and hot-water service within acceptable bounds, while preference constraints enforce availability windows, maximum interruption tolerance, and appliance completion requirements. Carbon constraints can either discourage discretionary consumption during high-carbon periods or prioritize load shifting toward cleaner intervals when comfort and service constraints permit. Dynamic carbon factor studies show that such signals can be incorporated as either objectives or operating limits, depending on the aggregator’s decarbonization policy [25].

Human-in-the-loop interaction mechanism

The human-in-the-loop mechanism operates as a review-and-reoptimization process. The model first produces a candidate dispatch schedule, then presents the operator with comfort-risk, carbon-risk, participation-burden, and feeder-impact indicators for each relevant action. The operator may lock, reject, or relax selected dispatch decisions before the model is solved again under the revised decision rules. This interaction is important because real demand response events can produce occupant discomfort or manual overrides when automated schedules fail to reflect local context [26].

Solution Method

Linearization and MILP conversion

The model is implemented as a mixed-integer linear program by representing device commitment, appliance start decisions, and participation states with binary variables, while comfort, power, and service states are represented with continuous variables. Thermal and service dynamics are approximated with linear relationships to maintain tractability and auditability for operator review. Preference costs are represented through piecewise-linear penalty structures so that willingness-to-accept can be included without requiring nonlinear optimization. Although reinforcement learning has been applied to demand response dispatch, a MILP is selected here because it provides transparent feasibility guarantees and explicit constraint interpretation [27].

Solver configuration

The dispatch process uses a rolling-horizon implementation in which updated weather, tariff, carbon, and participation information are incorporated before each operational run. Warm starts are used after operator edits so that the revised problem can reuse the structure of the previous candidate schedule. Comfort-sensitive loads are assigned high-priority feasibility treatment to prevent the optimizer from relying on discomfort as a routine source of flexibility. This configuration follows applied building optimization work in which operational usefulness depends on reliable solution quality, interpretability, and repeatable execution [28].

Computational performance

Computational performance is evaluated qualitatively through feasibility, repeatability, operator responsiveness, and consistency of dispatch recommendations across scenario runs. The model is considered operationally suitable when the solver returns feasible candidate schedules, when operator edits can be incorporated without restructuring the entire problem, and when the dispatch logic remains explainable to the aggregator. Binary appliance and air-conditioning scheduling studies support this modeling choice because device-level interpretability is often more valuable for implementation than purely heuristic speed [29]. The case study therefore reports comparative operational patterns rather than unsupported claims about exact solve times.

Case Study

System description and data sources

The case study represents a residential building cluster with HVAC systems, electric water heaters, deferrable appliances, and optional flexible charging. Baseline demand profiles, weather-dependent thermal behavior, tariff schedules, and carbon-intensity signals are combined to create a realistic dispatch environment. The system is designed to expose conflicts among feeder peak management, household comfort, prosumer willingness to shift, and low-carbon scheduling. Occupant-centric demand response evidence supports this cluster-level setting because comfort response and override risk differ across homes, schedules, and device types [30].

Scenario definitions

The evaluation compares a non-carbon-aware baseline, a fully automated carbon-aware dispatch, a preference-aware dispatch without human review, and the proposed human-in-the-loop carbon-aware dispatch. The baseline prioritizes cost and feeder operation but does not respond to carbon signals, while the automated carbon-aware case includes emissions information without operator intervention. The preference-aware case includes prosumer willingness-to-accept but does not allow post-optimization review. The proposed case combines carbon-aware dispatch, preference modeling, comfort protection, and operator review to assess the added value of human oversight [4].

Implementation details

All scenarios use the same device availability assumptions, weather conditions, tariff structure, and carbon-signal profile so that differences arise from the control logic rather than from inconsistent inputs. Sensitivity analysis varies comfort flexibility, carbon-policy strength, and participation availability in qualitative ranges rather than reporting unsupported empirical magnitudes. The model records the direction and relative strength of changes in cost, peak reduction, carbon savings, comfort violation, preference-regret, and operator intervention. Because enrollment and household response strongly influence realized flexibility, participation sensitivity is treated as a central part of the validation rather than a secondary robustness check [19].

Results and Discussion

Dispatch performance

The human-in-the-loop dispatch consistently shifts flexible consumption away from constrained feeder periods while avoiding excessive concentration of control actions on a small group of households. Compared with the non-carbon-aware baseline, the optimized schedules produce stronger peak-management behavior because HVAC cycling, water-heater interruption, and appliance deferral are coordinated rather than triggered independently. Compared with the fully automated carbon-aware schedule, the human-reviewed schedule is less aggressive in comfort-sensitive intervals but more balanced across participating prosumers. This trade-off is operationally important because demand response programs must preserve long-term participation rather than maximize short-term dispatch intensity [6].

Comfort and preference adherence

The proposed model produces better comfort and preference adherence than the fully automated carbon-aware case because operator review prevents the optimizer from repeatedly selecting high-inconvenience actions. Comfort-sensitive HVAC actions are reduced when they coincide with occupied periods, while lower-disruption appliance and water-heater flexibility is used more often where feasible. Preference-regret is also reduced because the model respects stated availability, load criticality, and participation limits. These results align with occupant-centered HVAC research showing that personalized comfort and override-aware control can maintain flexibility while reducing dissatisfaction [14].

Carbon savings

Carbon-aware dispatch shifts flexible load toward cleaner intervals when service windows and comfort constraints allow it. The largest carbon benefit occurs for deferrable appliances and water heating because these loads can often move across time without immediate comfort effects. HVAC carbon shifting is more constrained because low-carbon intervals may conflict with thermal comfort or occupied-period acceptability. This pattern is consistent with carbon-oriented demand response studies showing that emissions benefits depend on both signal quality and the temporal flexibility of the controlled load [21].

Human-in-the-loop impact

Human review improves the operational acceptability of the optimized schedule by identifying actions that are technically feasible but likely to create discomfort, dissatisfaction, or repeated participation burden. Operator edits mainly affect comfort-sensitive dispatch and high-burden households, while preserving much of the carbon-aware and peak-reduction structure produced by the optimizer. The resulting schedule is therefore not simply less automated; it is a revised dispatch plan in which algorithmic efficiency is balanced with contextual judgment.

Figure 2 summarizes how the proposed human-in-the-loop carbon-aware dispatch balances peak reduction, emissions alignment, comfort protection, preference adherence, and operator acceptability across the evaluated scenario logic.




Figure 2. Scenario-Based Evidence Map of Trade-Offs in Human-Reviewed Carbon-Aware Demand Response Dispatch

Table 3 summarizes the comparative scenario outcomes without introducing unsupported empirical magnitudes.

Table 3. Comparative Evaluation of Dispatch Scenarios Incorporating Cost, Peak, Carbon, Comfort, and Preference Outcomes with Human-in-the-Loop Review

Scenario	Cost outcome	Peak-management outcome	Carbon outcome	Comfort outcome	Preference outcome	Human-review role
Non-carbon-aware baseline	Reference operating case	Weakest peak coordination	Weakest carbon alignment	Best comfort preservation because little control is applied	No explicit preference optimization	No review
Cost-and-peak optimized, no carbon	Lower cost than baseline	Strong peak reduction	Carbon benefit depends on coincidence with cleaner periods	Moderate comfort burden	Moderate preference burden	No review
Fully automated carbon-aware dispatch	Slightly higher operating burden than cost-only dispatch	Strong peak reduction	Strong carbon alignment	Highest comfort burden among optimized cases	Higher preference burden due to aggressive automation	No review
Preference-aware dispatch without human review	Balanced operating burden	Moderate-to-strong peak reduction	Moderate-to-strong carbon alignment	Improved comfort relative to fully automated dispatch	Improved preference adherence	No review
Proposed human-in-the-loop dispatch	Balanced operating burden with	Strong but less aggressive peak reduction	Strong carbon alignment	Best comfort outcome among optimized cases	Best preference adherence	Operator locks, rejects, or

	operator-approved trade-offs		retained where feasible		among optimized cases	revises selected actions
--	------------------------------	--	-------------------------	--	-----------------------	--------------------------

Conclusion

This article formulated a human-in-the-loop optimization model for demand response dispatch under comfort constraints, prosumer preferences, and low-carbon flexibility signals. The framework coordinates flexible residential loads while preserving service quality and allowing operator review before implementation. The resulting model is intended as an operational decision-support tool rather than a purely conceptual architecture.

The case study demonstrates that human-reviewed carbon-aware dispatch can preserve much of the flexibility value of automated optimization while improving comfort and preference outcomes. The key advantage of the human-in-the-loop design is not that it replaces mathematical optimization, but that it constrains and corrects optimization when local context matters. This is especially important for comfort-sensitive loads and households with limited willingness to accept repeated dispatch.

The sensitivity analysis shows that comfort flexibility, carbon-policy strength, and participation availability shape the feasible demand response envelope. Wider comfort tolerance and stronger carbon signals increase the opportunity for load shifting, but only when household participation and service constraints remain acceptable. The framework therefore highlights the practical need to treat demand response as a socio-technical resource.

Future extensions should include real-time telemetry, adaptive preference learning, distribution-network power-flow constraints, and field deployment with actual operator feedback. Broader prosumer engagement could also improve parameter quality and reduce uncertainty in willingness-to-accept estimates. The proposed framework provides a foundation for demand response systems that are flexible, low-carbon, and accountable to human judgment.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Kim Y.J. A supervised-learning-based strategy for optimal demand response of an HVAC system in a multi-zone office building. *IEEE Trans Smart Grid*. 2020;11(5):4212-26.
- Tang R, Wang S. Model predictive control for thermal energy storage and thermal comfort optimization of building demand response in smart grids. *Appl Energy*. 2019;242:873-82.
- Jung W, Jazizadeh F. Human-in-the-loop HVAC operations: A quantitative review on occupancy, comfort, and energy-efficiency dimensions. *Appl Energy*. 2019;239:1471-508.
- Sarran L, Gunay HB, O'Brien W, Hviid CA, Rode C. A data-driven study of thermostat overrides during demand response events. *Energy Policy*. 2021;153:112290.
- Stenner K, Frederiks ER, Hobman EV, Cook S. Willingness to participate in direct load control: The role of consumer distrust. *Appl Energy*. 2017;189:76-88.
- Parrish B, Heptonstall P, Gross R, Sovacool BK. A systematic review of motivations, enablers and barriers for consumer engagement with residential demand response. *Energy Policy*. 2020;138:111221.
- Fleschutz M, Bohlayer M, Braun M, Henze G, Murphy MD. The effect of price-based demand response on carbon emissions in European electricity markets: The importance of adequate carbon prices. *Appl Energy*. 2021;295:117040.
- Zhang H, Xiao F, Zhang C, Li R. A multi-agent system based coordinated multi-objective optimal load scheduling strategy using marginal emission factors for building cluster demand response. *Energy Build*. 2023;281:112765.
- Mota B, Faria P, Vale Z. Residential load shifting in demand response events for bill reduction using a genetic algorithm. *Energy*. 2022;260:124978.
- Sridhar A, Honkapuro S, Ruiz F, Stoklasa J, Annala S, Wolff A, et al. Toward residential flexibility—Consumer willingness to enroll household loads in demand response. *Appl Energy*. 2023;342:121204.
- Ma S, Liu H, Wang N, Huang L, Su J, Zhao T. Incentive-based integrated demand response with multi-energy time-varying carbon emission factors. *Appl Energy*. 2024;359:122763.
- Chen Y, Luo F, Dong Z, Meng K, Ranzi G, Wong KP. A day-ahead scheduling framework for thermostatically controlled loads with thermal inertia and thermal comfort model. *J Mod Power Syst Clean Energy*. 2019;7(3):568-78.
- Nguyen AT, Ahn Y. Human-in-the-Loop Artificial Intelligence in the Energy Sector: A systematic review of paradigm shifts toward Industry 5.0. *Adv Appl Energy*. 2026:100267.
- Meimand M, Jazizadeh F. A personal touch to demand response: An occupant-centric control strategy for HVAC systems using personalized comfort models. *Energy Build*. 2024;303:113769.
- Frahm M, Dengiz T, Zwickel P, Maaß H, Matthes J, Hagenmeyer V. Occupant-oriented demand response with multi-zone thermal building control. *Appl Energy*. 2023;347:121454.

Kaspar K, Nweye K, Buscemi G, Capozzoli A, Nagy Z, Pinto G, et al. Effects of occupant thermostat preferences and override behavior on residential demand response in CityLearn. *Energy Build.* 2024;324:114830.

Dengiz T, Jochem P, Fichtner W. Demand response with heuristic control strategies for modulating heat pumps. *Appl Energy.* 2019;238:1346-60.

Sridhar A, Honkapuro S, Ruiz F, Stoklasa J, Annala S, Wolff A, et al. Residential consumer preferences to demand response: Analysis of different motivators to enroll in direct load control demand response. *Energy Policy.* 2023;173:113420.

Sridhar A, Honkapuro S, Ruiz F, Stoklasa J, Annala S, Wolff A. Residential consumer enrollment in demand response: An agent based approach. *Appl Energy.* 2024;374:123988.

Morstyn T, McCulloch MD. Multiclass energy management for peer-to-peer energy trading driven by prosumer preferences. *IEEE Trans Power Syst.* 2018;34(5):4005-14.

Wan T, Tao Y, Qiu J, Lai S. Internet data centers participating in electricity network transition considering carbon-oriented demand response. *Appl Energy.* 2023;329:120305.

Hlalele TG, Zhang J, Naidoo RM, Bansal RC. Multi-objective economic dispatch with residential demand response programme under renewable obligation. *Energy.* 2021;218:119473.

Zhao R, Song Z, Xu Y, Lu J, Guo W, Li H. Low-carbon demand response program for power systems considering uncertainty based on the data-driven distributionally robust chance constrained optimization. *IET Renew Power Gener.* 2025;19(1):e13021.

Zhou C, Zheng J, Liu S, Liu Y, Mei F, Pan Y, et al. Operation optimization of multi-district integrated energy system considering flexible demand response of electric and thermal loads. *Energies.* 2019;12(20):3831.

Zhang Q, Wang X, Ding J, Wang H, Zhao F, Ju X, et al. A Framework for sustainable power demand response: Optimization scheduling with dynamic carbon emission factors and dual DPMM-LSTM. *Sustainability.* 2025;17(20):9123.

Le Cam A, Southernwood J, Ring D, Clarke D, Creedon R. Impact of demand response on occupants' thermal comfort in a leisure center. *Energy Eff.* 2021;14(8):91.

Bahrami S, Chen YC, Wong VW. Deep reinforcement learning for demand response in distribution networks. *IEEE Trans Smart Grid.* 2020;12(2):1496-506.

Meinrenken CJ, Mehmani A. Concurrent optimization of thermal and electric storage in commercial buildings to reduce operating cost and demand peaks under time-of-use tariffs. *Appl Energy.* 2019;254:113630.

Bao YQ, Pei MX, Wu XH, Yu ZC. Price-based optimization of air conditioners by combined temperature set-point and main-power ON/OFF control. *Energy Build.* 2023;296:113335.

Kini RL, Vlachokostas A, Brambley MR, Rogers A. Occupant-Centric Demand Response for Thermostatically-Controlled Home Loads. *IEEE Trans Smart Grid.* 2025;16(3):2234-45.